arXiv:1602.03865v1 [cs.CG] 11 Feb 2016

HIGHER SIGNATURE DELAUNAY DECOMPOSITIONS

JEFFREY DANCIGER, SARA MALONI, AND JEAN-MARC SCHLENKER

Abstract. A Delaunay decomposition is a cell decomposition in $\mathbb{R}^d$ for which each cell is inscribed in a Euclidean ball which is empty of all other vertices. This article introduces a generalization of the Delaunay decomposition in which the Euclidean balls in the empty ball condition are replaced by other families of regions bounded by certain quadratic hypersurfaces. This generalized notion is adaptable to geometric contexts in which the natural space from which the point set is sampled is not Euclidean, but rather some other flat semi-Riemannian geometry, possibly with degenerate directions. We prove the existence and uniqueness of the decomposition and discuss some of its basic properties. In the case of dimension $d = 2$, we study the extent to which some of the well-known optimality properties of the Euclidean Delaunay triangulation generalize to the higher signature setting. In particular, we describe a higher signature generalization of a well-known description of Delaunay decompositions in terms of the intersection angles between the circumscribed circles.

Contents

1. Introduction 2
2. Classical Delaunay decomposition and its generalizations 3
3. The geometry associated to a quadratic form 7
3.1. Non-degenerate quadratic forms 7
3.2. Degenerate quadratic forms 9
3.3. The convex hull construction revisited 11
3.4. Möbius invariance of Q–Delaunay decompositions 13
4. Higher signature Delaunay decompositions in dimension $d = 2$ 16
4.1. Interior angles and edge angles of higher signature Delaunay triangulations 16
4.2. Interior angles and edge angles 17
4.3. Prescribing the edge angles 18
4.4. Interior angle optimization 19
5. Outline of applications 21
5.1. Triangulations in Minkowski space 21
5.2. Analysis of discrete functions 21
5.3. Proximity Graphs 22
5.4. Minimizing interpolation error 23
5.5. Other applications 24
5.6. Implementation 24
References 24

Date: February 12, 2016(v0).

Danciger was partially supported by grant DMS-1510254 and Maloni was partially supported by grant DMS-1506920, from the National Science Foundation. The authors also acknowledge support from U.S. National Science Foundation grants DMS-1107452, 1107263, 1107367 "RNMS: GEometric structures And Representation varieties" (the GEAR Network).

1. Introduction

Let X be a finite set of points in Euclidean space $\mathbb{R}^d$ of dimension $d \geq 2$. A cell decomposition of the convex hull $\mathrm{CH}(X)$ of X is called a *Delaunay decomposition* for X if it satisfies the *empty ball condition:* each cell of the decomposition is inscribed in a unique round ball which does not contain any points of X in its interior. A classical theorem of Delaunay [8] states that if X is in generic position (meaning that no $d+1$ points lie in an affine hyperplane and no $d+2$ points lie on any sphere), then there exists a unique Delaunay decomposition for X, all of the cells of which are simplices.

This article introduces a generalization of the Delaunay decomposition in which the balls in the empty ball condition above are replaced by other families of regions bounded by certain quadratic hypersurfaces. This generalized notion of Delaunay decomposition is adaptable to geometric contexts in which the natural space from which the point set X is sampled is not Euclidean, but rather some other flat semi-Riemannian geometry, possibly with degenerate directions. For example, in the Minkowski plane $\mathbb{R}^{1,1}$, we replace Euclidean circles with the hyperbolas of constant Minkowski distance from a point. In the degenerate plane $\mathbb{R}^{1,0,1}$, thought of as a rescaled limit of the Euclidean plane in which one direction has become infinitesimal, we replace Euclidean circles with their rescaled limits, namely parabolas.

A quadratic form Q on $\mathbb{R}^d$ defines a natural geometric structure on $\mathbb{R}^d$. In the case that Q is non-degenerate, we simply use Q in place of the Euclidean norm to measure distances (possibly imaginary), angles, and other geometric quantities. The points within a constant Q-distance of a fixed point are called Q-balls and we substitute these in place of Euclidean balls to define a natural notion of a Q–Delaunay triangulation. In the case that Q has some degenerate directions, we study a natural geometry which treats the degenerate directions as infinitesimal and define an appropriate notion of Q-ball which is slightly more subtle than in the non-degenerate case. Again we define a Delaunay decomposition in terms of an appropriate empty ball condition. We stress that in the case that Q is not positive (or negative) definite, Q-Delaunay decompositions are not in general obtained via projective transformations from Euclidean ones.

In Section 2 we extend a classical proof of the existence and uniqueness of Delaunay decompositions in the Euclidean case to the current setting under the necessary assumption that the point set is in *spacelike position* with respect to Q, meaning that any two points are positive real distance apart.

Theorem A. *Let Q be a quadratic form on $\mathbb{R}^d$ and let $X \subset \mathbb{R}^d$ be a finite set in spacelike position with respect to Q and in generic position (meaning that no $d+1$ points lie in an affine hyperplane and no $d+2$ points lie on any Q–sphere). Then there exists a unique Delaunay decomposition for X. We will denote this decomposition by* $\mathrm{Del}_Q(X)$

In Section 3 we will review some details on quadratic forms on $\mathbb{R}^d$ and their associated geometry. In the case that the quadratic form Q is not degenerate, $\mathrm{Del}_Q(X)$ is naturally dual to a generalized Voronoi decomposition, the vertices of which are the centers of the empty Q–balls of $\mathrm{Del}_Q(X)$ (Proposition 3.3). The case that Q is degenerate is more delicate. In this case it is natural to split $\mathbb{R}^d = \mathbb{R}^m \oplus \mathbb{R}^{d-m}$ as the sum of a non-degenerate subspace $\mathbb{R}^m$ complementary to the degenerate subspace $\mathbb{R}^{d-m}$. The associated geometry that we study treats the degenerate directions as infinitesimal so that a point (p_0, v) in this geometry is thought as a point p_0 in the non-degenerate subspace $\mathbb{R}^m$ plus infinitesimal information $v \in \mathbb{R}^{d-m}$ describing the normal derivative of a path p_t leaving the subspace $\mathbb{R}^m$ going into $\mathbb{R}^d$. We show in Theorem 3.6 that the Delaunay decomposition $\mathrm{Del}_Q(X)$ in this case predicts

for short time $t > 0$ the combinatorics of the path of Delaunay decompositions $\mathrm{Del}_{Q'}(X_t)$ for points X_t whose first order data agrees with that determined by X and with respect to a non-degenerate quadratic form Q' obtained by making the infinitesimal directions finite.

In Section 4, we study the extent to which some of the well-known optimality properties of the Euclidean Delaunay triangulation generalize to the non-Euclidean setting in the case of dimension $d = 2$. There are essentially two non-Euclidean quadratic forms Q, one Lorentzian (signature $(1, 1)$) and one degenerate. In the first case, the Q-circles are hyperbolas and in the second case they are parabolas. As in the Euclidean case, in either of these geometries weights may be assigned to the edges of a triangulation by summing the angles at opposite vertices, or alternatively by measuring the angle of intersection of the circumscribed Q–circles containing a given edge. Using recent results [7] about three-dimensional ideal polyhedra in constant negative curvature geometries, in Theorem 4.7 we characterize the possible edge weights that may occur for a Q–Delaunay triangulation in either case. In addition, using a generalized Thales' Theorem, Theorem 4.8 proves that Q–Delaunay triangulations are optimally fat: they maximize the angle sequence (ordered lexicographically) over all triangulations as in the Euclidean case.

In Section 5, we briefly outline some possible applications of the higher signature Delaunay decompositions developed here, and mention a few open questions.

2. Classical Delaunay decomposition and its generalizations

Recall that a ball is defined in terms of the Euclidean norm, $\|x\|^2 = x^T x$. Specifically, the ball of radius-squared $D \in \mathbb{R}$ and center $p \in \mathbb{R}^d$ is the set of points $x \in \mathbb{R}^d$ satisfying the inequality:

$$\|x - p\|^2 \leq D. \tag{1}$$

This may be expressed as the equivalent condition that:

$$\|x\|^2 \leq \varphi(x) + D', \tag{2}$$

where $\varphi : \mathbb{R}^d \to \mathbb{R}$ is the linear functional defined by $\varphi(x) = 2p^T x$ and the constant $D' = D - \|p\|^2$. In other words, a ball is exactly the set of points in $\mathbb{R}^d$ whose image in $\mathbb{R}^{d+1}$ on the graph of the function $\|\cdot\|^2$ lies below some affine hyperplane. The norm-squared function $x \mapsto \|x\|^2 = x^T x$ is one example of a *quadratic form*, which we define here as any function $Q : \mathbb{R}^d \to \mathbb{R}$ of the form

$$Q(x) = x^T A x,$$

where A is some symmetric $d \times d$ matrix. If the matrix A is non-singular, then we call Q *non-degenerate*, and if the matrix A is positive definite (respectively indefinite), then we call Q *positive definite* (respectively *indefinite*). Up to transforming $\mathbb{R}^d$ by a linear isomorphism, a quadratic form Q is determined entirely by the number of positive, negative, and zero eigenvalues of A. This data is referred to as the signature of Q.

Let Q be a non-degenerate quadratic form, possibly indefinite, on $\mathbb{R}^d$. The term *Q–ball* refers to any region defined by the inequality

$$Q(x - p) \leq D, \tag{3}$$

where $p \in \mathbb{R}^d$ and $D \in \mathbb{R}$ (possibly negative). Equivalently, a Q–ball is any region defined by the inequality

$$Q(x) \leq \varphi(x) + D', \tag{4}$$

where $\varphi : \mathbb{R}^d \to \mathbb{R}$ is any linear functional and $D' \in \mathbb{R}$ is any constant.

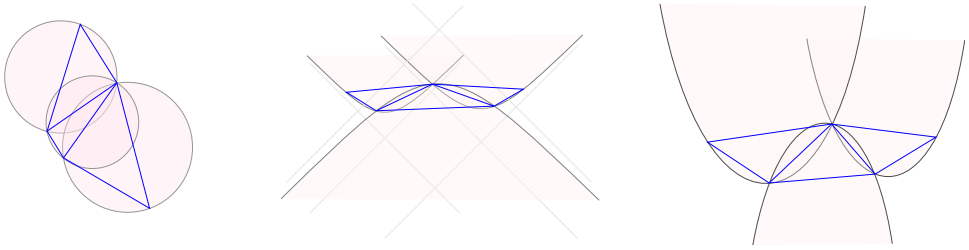

FIGURE 1. An example of the Q–Delaunay decomposition of five points in spacelike general position in the case that:

- (left) $Q(x_1, x_2) = x_1^2 + x_2^2$ is the standard Euclidean norm, and so the Q–balls are bounded by circles;
- (middle) $Q(x_1, x_2) = x_1^2 - x_2^2$ is the standard Minkowski norm, and so the Q–balls are bounded by hyperbolas;
- (right) $Q(x_1, x_2) = x_1^2$ is a degenerate norm, in which case the Q–balls are bounded by parabolas.

Of course, Inequality (3) may be described in the form (4) by simply taking $\varphi(x) = 2p^T Ax$ and $D' = D + Q(p)$. However, solving for p, D in terms of φ, D' in order to deduce inequality (3) from (4) requires the quadratic form Q (that is, the matrix A) to be non-degenerate. Hence, Inequalities (3) and (4) define different families of regions in the case that Q is degenerate. It turns out that the regions defined by Inequality (4) define a much more interesting theory of Delaunay decomposition in this case.

Definition 2.1 (Q–ball)**.** Let Q be any quadratic form (possibly degenerate) on $\mathbb{R}^d$. The term *Q–ball* refers to the region defined by Inequality (4), where $\varphi\colon \mathbb{R}^d \to \mathbb{R}$ is a linear functional and $D' \in \mathbb{R}$ is a constant. The boundary of a Q–ball is called a *Q–sphere*.

We will see in Section 3.2 that a degenerate form Q defines a special geometry in which the degenerate directions of Q are treated as infinitesimal and we think of Q as a degeneration of a family of non-degenerate quadratic forms Q'_t as $t \to 0$. Under that interpretation, Q–balls may be thought of as limits of Q'_t-balls with center escaping to infinity as $t \to 0$.

In order to obtain a sensible theory of Delaunay decompositions using a semi-definite quadratic form Q, we must restrict the point sets under consideration. We make the following definition:

Definition 2.2. Let Q be a quadratic on $\mathbb{R}^d$ (possibly indefinite and/or degenerate). A finite set of points X in $\mathbb{R}^d$ is said to be in *spacelike position* with respect to Q if, for any distinct points $x, y \in X$, the positivity condition $Q(x - y) > 0$ holds.

Replacing the notion of ball with that of Q–ball in the definition of the classical Delaunay decomposition yields the following generalization.

Definition 2.3 (Q–Delaunay decomposition)**.** Let Q be any quadratic form on $\mathbb{R}^d$ and let X be a finite set of points in $\mathbb{R}^d$ in spacelike position with respect to Q. A cell decomposition of the convex hull $\mathrm{CH}(X)$ of X with 0-skeleton X is called a *Q–Delaunay decomposition* if it satisfies the *empty Q–ball condition:* each cell of the decomposition is inscribed in a unique Q–ball which does not contain any points of X in its interior. We will call the boundary of such a Q–ball an *empty Q–sphere*.

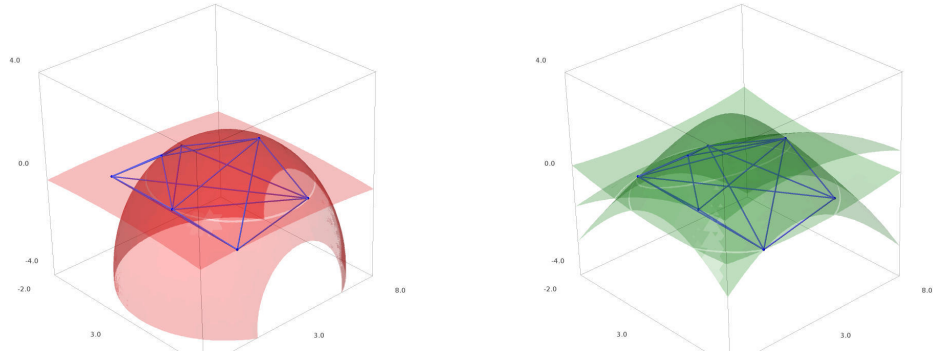

FIGURE 2. The Q-Delaunay decompositions of the same 7 points when Q is the Euclidean (left) and the Minkowski (right) bilinear form on $\mathbb{R}^3$. Only some of the Q-balls have been drawn.

We now prove Theorem A. The proof is adapted from a well-known proof in the case that Q is the standard Euclidean norm, see [2].

Proof of Theorem A. Given the quadratic form Q on $\mathbb{R}^d$, we consider the quadratic form $\widehat{Q}$ on $\mathbb{R}^{d+2}$ defined by

$$\widehat{Q}(x_1, \ldots, x_d, x_{d+1}, x_{d+2}) := Q(x_1, \ldots, x_d) - x_{d+1}x_{d+2}.$$

Associate to $\widehat{Q}$, we define the open subset $\mathbb{X} \subset \mathbb{RP}^{d+1}$ as follows:

$$\mathbb{X} = \{x \in \mathbb{R}^{d+2} \setminus \{0\} \mid \widehat{Q}(x) < 0\}/\mathbb{R}^*,$$

that is the set of negative lines in $\mathbb{R}^{d+2}$ with respect to $\widehat{Q}$. Then $\mathbb{X}$ is a model for semi-Riemannian geometry, possibly with degenerate directions, associated with $\widehat{Q}$ with constant negative sectional curvature. In the affine chart $x_{d+2} = 1$, the ideal boundary of $\mathbb{X}$ is described by the equation

$$x_{d+1} = Q(x_1, \ldots, x_d),$$

i.e. the graph $\mathscr{P}$ in $\mathbb{R}^{d+1}$ of Q. The projection map $\varpi : \mathbb{R}^{d+1} \to \mathbb{R}^d$ onto the first d-coordinates, defined by

$$\varpi(x, x_{d+1}) = x,$$

restricts to a homeomorphism $\mathscr{P} \to \mathbb{R}^d$, which, in the case that Q is the usual Euclidean norm-squared, is usually called *stereographic projection*.

Now, consider the convex hull $\mathscr{C}$ of the inverse images $X' = \varpi^{-1}(X)$ of X on $\mathscr{P}$. $\mathscr{C}$ may be regarded as a convex polyhedron in $\mathbb{X}$ whose vertices lie on the ideal boundary $\partial\mathbb{X}$. By the genericity assumption, $\mathscr{C}$ is a $d+1$-dimensional polyhedron, whose boundary is a union of d-dimensional simplices which come in two types, called *bottom* and *top*. A face f on the boundary $\partial\mathscr{C}$ of $\mathscr{C}$ is a bottom face if its unit normal, directed toward the interior of $\mathscr{C}$, has positive x_{d+1} component. A face f is a top face if its inward directed unit normal has negative x_{d+1} component. The unit normal to a face never has zero x_{d+1} component by the genericity assumption.

We claim that ϖ restricts to a homeomorphism from the union $\partial\mathscr{C}_-$ of the bottom faces of $\partial\mathscr{C}$ to the convex hull $\mathrm{CH}(X)$ of X in $\mathbb{R}^d$ and that ϖ takes the faces of $\partial\mathscr{C}_-$ to the cells of the unique Q–Delaunay decomposition for X. To see this, observe that the condition that all points $(x, Q(x))$ of X' lie (weakly) above

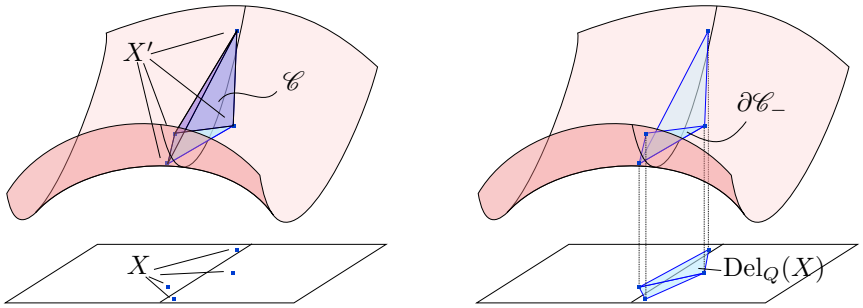

FIGURE 3. To find the Q–Delaunay decomposition $\mathrm{Del}_Q(X)$, first lift X to its image X' on the graph of Q and take the convex hull $\mathscr{C}$ of X (left panel). Then project the bottom faces $\partial\mathscr{C}_-$ down to $\mathbb{R}^d$ (right panel).

the affine hyperplane spanned by some face f of $\partial\mathscr{C}_-$ is described by an inequality of the form

$$Q(x) \geq \varphi(x) + D'$$

for some linear functional φ and constant D', with equality exactly when $(x, Q(x))$ is a vertex of f. Hence, the cell decomposition of $\mathbb{R}^d$ defined by projection of the cell structure of $\partial\mathscr{C}_-$ satisfies that the vertices of any face $\varpi(f)$ lie on a Q–sphere S_f which bounds a Q–ball B_f whose interior does not contain points of X. In the case that Q is positive definite the proof, which up to this point is well-known, is complete. There is, however, one subtle point to check if Q is not positive definite, which is the case of interest in this article. Definition 2.3 requires that the simplex $\varpi(f)$ be *inscribed* in B_f: so we must check that $\varpi(f)$ is actually contained in B_f. This will follow from the assumption that X is in spacelike position. Let $\langle\cdot,\cdot\rangle_Q$ denote the inner product associated to Q, defined by $\langle x, y\rangle_Q = x^T A y$ for all $x, y \in \mathbb{R}^d$. Then

$$Q(x-y) > 0 \iff 2\langle x, y\rangle_Q < Q(x) + Q(y). \tag{5}$$

Let $x'_1, \ldots, x'_m$ be vertices of the face f, where $x'_k = (x_k, Q(x_k))$, for each $k \in \{1, \ldots, m\}$. A typical point of f is a convex combination of the form $y = t_1 x_1 + \cdots + t_m x_m$, where $t_1, \ldots, t_m$ are positive coefficients such that $t_1 + \cdots + t_m = 1$. Then

$$\begin{aligned}
Q(y) &= \left\langle \sum_{i=1}^m t_i x_i, \sum_{j=1}^m t_j x_j \right\rangle_Q \\
&= \sum_{i=1}^m t_i^2 Q(x_i) + 2\sum_{j<k} t_j t_k \langle x_j, x_k\rangle_Q \\
&\leq \sum_{i=1}^m t_i^2 Q(x_i) + \sum_{j<k} t_j t_k (Q(x_j) + Q(x_k)) \\
&= \sum_{i=1}^m t_i (t_1 + \cdots + t_m) Q(x_i) = \sum_{i=1}^m t_i Q(x_i) \\
&= \sum_{i=1}^m t_i (\varphi(x_i) + D) = \varphi(y) + D \\
&\implies Q(y) \leq \varphi(y) + D.
\end{aligned}$$

The hypotheses that the points $x_1, \dots, x_m$ are in spacelike position was applied (via (5)) in the third line. Note that equality holds if and only if $y = x_k$ for some $k \in \{1, \dots, m\}$. It follows that $\varpi(f) \subset B_f$ with $\varpi(f) \cap S_f = \{x_1, \dots, x_m\}$, that is $\varpi(f)$ is inscribed in B_f. The same ideas can be used to prove the uniqueness. □

Remark 2.4. Note that ϖ restricts to a homeomorphism from the union $\partial\mathscr{C}_+$ of the top faces of $\partial\mathscr{C}$ to the convex hull $\mathrm{CH}(X)$ of X in $\mathbb{R}^d$ and that ϖ takes the faces of $\partial\mathscr{C}_+$ to the cells of the unique decomposition for X where each Q–ball satisfies the *full Q–ball condition*: each cell of the decomposition is inscribed in a unique Q–ball such that it contains all the points of X in its closure. We will call the boundary of such a Q–ball a *full Q–sphere*.

Remark 2.5. In the familiar case that Q is the Euclidean norm, the surface $\mathscr{P}$ in the proof above may be naturally interpreted as the ideal boundary of the hyperbolic space $\mathbb{H}^{d+1}$ with one point removed. The convex hull $\mathscr{C}$ of X' is an ideal polyhedron in $\mathbb{H}^{d+1}$, the geometry of which gives information about the geometry of the Delaunay decomposition of X. In the general case a similar interpretation is possible. If Q is non-degenerate of signature (p, q), then $\mathscr{P}$ identifies with (a dense subset of) the ideal boundary of the negatively curved semi-Riemannian space $\mathbb{H}^{p+1,q}$ of signature $(p + 1, q)$ and constant negative curvature. If Q is degenerate, then $\mathscr{P}$ may be interpreted as (a dense subset of) the boundary of a degenerate geometry obtained from some $\mathbb{H}^{p+1,q}$ as a limit. See Section 3.3.

3. The geometry associated to a quadratic form

3.1. Non-degenerate quadratic forms.

3.1.1. *Flat semi-Riemannian geometry.* Throughout this section, we assume Q is a non-degenerate quadratic form on $\mathbb{R}^d$. The numbers p and q of positive and negative eigenvalues of the matrix A of Q satisfy that $p + q = d$. The pair (p, q) is called the signature of Q. After changing coordinates by a linear transformation, we may assume that Q is the standard quadratic form of signature (p, q), meaning its matrix $A = I_p \oplus -I_q$ is simply the diagonal matrix with p $(+1)$'s followed by q (-1)'s on the diagonal. The inner product $\langle \cdot, \cdot \rangle_Q$ induced by Q is given by the simple expression

$$\langle x, y \rangle_Q = x_1 y_1 + \cdots + x_p y_p - x_{p+1} y_{p+1} - \cdots - x_d y_d \ .$$

We assume that $p, q > 0$, since the case $q = 0$ is the classical Euclidean case and the case $p = 0$ is un-interesting (there are no point sets X in spacelike position when $|X| > 1$).

The inner product $\langle \cdot, \cdot \rangle_Q$ defines, by parallel translation, a flat semi-Riemannian metric on the affine space of $\mathbb{R}^d$. This space and metric together is denoted by $\mathbb{R}^{p,q}$. It is the simply connected model for flat semi-Riemannian geometry of signature (p, q). The isometry group of $\mathbb{R}^{p,q}$ is precisely

$$\mathrm{Isom}(\mathbb{R}^{p,q}) = \mathrm{Aff}(\mathrm{O}(Q)) = \mathrm{O}(p, q) \ltimes \mathbb{R}^d,$$

where $\mathrm{O}(Q) = \mathrm{O}(p, q)$ denotes the linear transformations of $\mathbb{R}^d$ preserving Q. Of course, any natural construction in flat semi-Riemannian geometry of signature (p, q) should be invariant under this group.

Proposition 3.1. *The set of Q–balls is invariant under* $\mathrm{Isom}(\mathbb{R}^{p,q})$*, hence so is the notion of Q–Delaunay decomposition: if $f \in \mathrm{Isom}(\mathbb{R}^{p,q})$ and $X \subset \mathbb{R}^{p,q}$ is in spacelike and generic position, then* $\mathrm{Del}_Q(f(X)) = f(\mathrm{Del}_Q(X))$.

In Section 3.4, we will show that in fact the Q–Delaunay decomposition is invariant under a larger set of transformations, specifically a natural subset (depending on X) of Möbius transformations (Theorem 3.15).

The quadratic form Q may be used to measure lengths, just as in the Euclidean setting, with some minor differences. First, note that the Q–distance-squared $Q(x-y) = \langle x-y, x-y\rangle_Q$ between two distinct points $x, y \in \mathbb{R}^d$ may be positive, negative, or zero. Borrowing terminology from the Lorentzian setting ($p = d-1, q = 1$), we call the displacement between x and y *spacelike* if $Q(x-y) > 0$, *timelike* if $Q(x-y) < 0$, or *lightlike* if $Q(x-y) = 0$. When $Q(x-y) < 0$, one may think of the distance between x and y as an imaginary number, however it is more convenient to simply work with $Q(x-y)$ rather than its square root.

3.1.2. *Q–Voronoi decompositions.* As in the classical setting, the Q–Delaunay decomposition is dual to an analogue of the Voronoi decomposition, which we call the Q–Voronoi decomposition.

Definition 3.2. Let $X \subset \mathbb{R}^d$ be finite and let Q be a non-degenerate quadratic form. Then for each $x \in X$, we define the Voronoi region of x to be

$$\mathrm{Vor}_Q(x) := \{y \in \mathbb{R}^d : Q(x-y) \leq Q(x'-y) \text{ for all } x' \in X\}.$$

Proposition 3.3. *Suppose $X \subset \mathbb{R}^d$ is a finite set in spacelike and generic position with respect to Q. Then*

1. *For each $x \in X$, $\mathrm{Vor}_Q(x)$ is a (possibly unbounded) convex polyhedron, with timelike codimension 1 faces, containing x in its interior.*
2. *The collection of $\mathrm{Vor}_Q(x)$ for all $x \in X$ gives a polyhedral decomposition of $\mathbb{R}^d$ which we call the Q–*Voronoi decomposition *with respect to X. It will be denoted by $\mathrm{Vor}_Q(X)$.*
3. *For each facet f of codimension k of $\mathrm{Vor}_Q(X)$, there exists $k+1$ unique points $p_0, p_1, \ldots, p_k$, distinct from one another, such that*
$$f = \mathrm{Vor}_Q(p_0) \cap \mathrm{Vor}_Q(p_1) \cap \cdots \cap \mathrm{Vor}_Q(p_k).$$
4. *The vertices of $\mathrm{Vor}_Q(X)$ are the centers of the empty Q–balls of $\mathrm{Del}_Q(X)$.*

Remark 3.4. It is also natural to define the *inverse Voronoi region* as follows: for each $x \in X$, consider the set

$$\mathrm{Vor}_Q^{-1}(x) = \{y \in \mathbb{R}^d : Q(x-y) \geq Q(x'-y) \text{ for all } x' \in X\}.$$

Then, denoting by $\mathrm{Vor}_Q^{-1}(X)$ the collection of $\mathrm{Vor}_Q^{-1}(x)$ for all $x \in X$, the vertices of $\mathrm{Vor}_Q^{-1}(X)$ are the centers of the full Q–balls, as defined in Remark 2.4.

Proof. Since X is in spacelike position, for all distinct $x, p \in X$, the set of points $y \in \mathbb{R}^d$ such that $Q(x-y) \leq Q(p-y)$ is a half-space, bounded by a timelike hyperplane, containing x in its interior. Therefore, $\mathrm{Vor}_Q(x)$ is the union of finitely many half-spaces (each bounded by a timelike hyperplane) and it is therefore a convex polyhedron (which might however be unbounded). Moreover $\mathrm{Vor}_Q(x)$ contains x in its interior, and the first point is proved.

The second point is clear since, by definition, each point $y \in \mathbb{R}^d$ is contained in at least one of the $\mathrm{Vor}_Q(x)$ for some $x \in X$.

The third point follows from the first point and from the hypothesis that X is in generic position: given a codimension k facet f of $\mathrm{Vor}_Q(X)$ in the boundary of $\mathrm{Vor}_Q(p_0)$, for some $p_0 \in X$, there are k other pairwise distinct points $p_1, \cdots, p_k \in X \setminus \{p_0\}$ such that for all $y \in f$ and for all $j \in \{0, \cdots, k\}$, we have $Q(y-p_i) = Q(p'_j - y)$. Moreover there cannot be more than k such points, by the genericity hypothesis.

The last point follows from the third point because a vertex v is a codimension d facet. Hence, v is Q-equidistant to $d+1$ points of X and is further away from the other points. □

Remark 3.5. Note that it is necessary to suppose that X is in spacelike position to have $x \in \mathrm{Vor}_Q(x)$. The hypothesis that X is in generic position, however, is not as essential. If it is not satisfied, the same convex hull construction naturally defines a Q–Delaunay decomposition some of whose cells may not be simplices. The codimension 0 non-simplicial cells of the decomposition then correspond to empty spheres containing more than $d+1$ points of X. The same phenomenon appears for the usual Delaunay decomposition in Euclidean space.

3.2. **Degenerate quadratic forms.** Consider a time-dependent finite point set X_t in $\mathbb{R}^d$. That is, for each $t \geq 0$, $X_t = \{p_1(t), \ldots, p_n(t)\}$, where for each $k \in \{1, \ldots, n\}$, $p_k(t)$ is a smooth path. We assume that at time $t = 0$, all of the points lie in a proper subspace $V \subset \mathbb{R}^d$ and are in generic position within that subspace. We are interested in the time-dependent Delaunay decomposition of X_t, with respect to some non-degenerate quadratic form Q on $\mathbb{R}^d$ (perhaps the standard Euclidean norm, or perhaps an indefinite form defining a flat semi-Riemannian structure as in the previous section). Of course, at time $t = 0$, the points X_0 are not in generic position and the Q-Delaunay decomposition is not defined (degenerate). However, we wish to use infinitesimal data at time $t = 0$ to predict the combinatorics of the Delaunay decomposition for $t > 0$.

Assume henceforth that the subspace V which contains the points at time $t = 0$ is non-degenerate with respect to Q, and let U denote the Q-orthogonal complement. Then, after changing coordinates appropriately, we may assume that $V = \mathbb{R}^m$ is the coordinate hyperplane corresponding to the first m coordinates and that $U = \mathbb{R}^{d-m}$ is the coordinate hyperplane corresponding to the remaining $d-m$ coordinates, so that the direct sum decomposition $\mathbb{R}^d = \mathbb{R}^m \oplus \mathbb{R}^{d-m}$ is Q–orthogonal. [1] In these coordinates we write $p_k(t) = (y_k(t), z_k(t))$ for all $1 \leq k \leq n$. Consider the set

1) S: Is it ok? JD: modified

$$X = \{q_k = (y_k(0), z_k'(0)) \mid 1 \leq k \leq n\},$$

where $z_k'(0)$ denotes the derivative of $z_k(t)$ at $t = 0$. Let $\mathscr{Q}$ denote the degenerate quadratic form defined by

$$\mathscr{Q}(y, z) = Q(y, 0) .$$

Theorem 3.6. *With notation as above, assume that X is in spacelike, generic position for $\mathscr{Q}$. Then for all $t > 0$ sufficiently small, the natural map $X_t \to X$ induces a cell-wise homeomorphism taking $\mathrm{Del}_Q(X_t)$ to $\mathrm{Del}_{\mathscr{Q}}(X)$. In other words, for short time, the combinatorics of the Q–Delaunay decomposition of X_t agrees with the combinatorics of the $\mathscr{Q}$-Delaunay decomposition of X.*

The theorem will follow from the next lemma which gives a natural interpretation of $\mathscr{Q}$-balls as rescaled limits of Q–balls under the rescaling maps $L_t : \mathbb{R}^d \to \mathbb{R}^d$ defined by

$$L_t(y, z) = \left(y, \frac{z}{t}\right).$$

The proof, omitted, is a simple computation.

Lemma 3.7. *Let B denote the $\mathscr{Q}$-ball defined by the equation $\mathscr{Q}(y, z) \leq \varphi(y, z) + D$ for some linear functional φ and constant D. Let B_t be a family of Q–balls, defined for $t > 0$, by the equations $Q(y, z) \leq \varphi_t(y, z) + D_t$ where*

$$\lim_{t \to 0} \varphi_t \circ L_t^{-1} = \varphi,$$
$$\lim_{t \to 0} D_t = D.$$

Then $L_t B_t \to B$ as $t \to 0$.

And now the proof of Theorem 3.6.

Proof of Theorem 3.6. Let I denote the $d+1$ indices of the points of X lying on the boundary of an empty $\mathscr{Q}$-ball

$$B^I = \{(y,z) \mid \varphi^I(y_i,z_i) + D^I\}$$

of $\mathrm{Del}_{\mathscr{Q}}(X)$, where the (unique) linear functional φ^I and constant D^I are defined by

$$\mathscr{Q}(y_i,z_i') = \varphi^I(y_i,z_i') + D^I \ \ \forall i \in I.$$

For every sufficiently small $t>0$, let φ_t^I and D_t^I be the unique linear functional and constant defining the Q–ball B_t^I containing $(y_i(t),z_i(t))$ in its boundary for all $i \in I$. Then

$$Q(y_i(t),z_i(t)) = \varphi_t^I(y_i(t),z_i(t)) + D_t^I$$

and therefore

$$\begin{aligned} Q\left(y_i(t), t\frac{z_i(t)}{t}\right) &= \varphi_t^I\left(y_i(t), t\frac{z_i(t)}{t}\right) + D_t^I \\ &= \varphi_t^I \circ L_t^{-1}\left(y_i(t), \frac{z_i(t)}{t}\right) + D_t^I \ . \end{aligned}$$

Taking the limit as $t \to 0$ we find that

$$Q(y_i(0),0) = \lim_{t\to 0}\left(\varphi_t^I \circ L_t^{-1}\right)(y_i(0),z_i'(0)) + \lim_{t\to 0} D_t^I \ ,$$

so that

$$\mathscr{Q}(y_i(0),z_i'(0)) = \lim_{t\to 0}\left(\varphi_t^I \circ L_t^{-1}\right)(y_i(0),z_i'(0)) + \lim_{t\to 0} D_t^I$$

for all $i \in I$. It now follows that $\varphi_t^I \circ L_t^{-1} \to \varphi$ and $D_t^I \to D^I$. We now apply Lemma 3.7 to obtain that $L_t B_t^I \to B^I$ as $t \to 0$. Observing that $L_t X_t \to X$, it now follows that for sufficiently small $t>0$, the rescaled balls $L_t B_t^I$ are empty of the other points of the rescaled point set $L_t X_t$. Hence, for sufficiently small $t>0$, the $\mathscr{Q}$–balls B_t^I are the empty balls defining the $\mathscr{Q}$–Delaunay triangulation of X_t. □

3.2.1. *The geometry of a degenerate quadratic form.* Let $\mathscr{Q}$ be a degenerate quadratic form as in the previous subsection. Parallel translation of the inner product $\langle\cdot,\cdot\rangle_{\mathscr{Q}}$ determines a flat semi-Riemannian metric on $\mathbb{R}^d$, which in this case is degenerate. The isometries of such a metric form an infinite-dimensional group, indicating that this degenerate geometry does not have very much structure. We may impose more structure by thinking of the degenerate directions as having infinitesimal length rather than zero length. Such structure is best described by a group G acting on $\mathbb{R}^d$ (the "symmetries of the structure") rather than a metric. We define G as a limit of the isometry groups of non-degenerate quadratic forms Q_t which are degenerating to $\mathscr{Q}$. Specifically, for each $t>0$, define Q_t by $Q_t(y,z) = Q(y,0) + tQ(0,z)$, and let $G_t = \mathrm{Isom}(Q_t)$. Define G to be the limit of G_t as $t \to 0$ in the Chabauty topology on the space of subgroups of the Lie group $\mathrm{Aff}(\mathbb{R}^d)$ of affine transformations of $\mathbb{R}^d$: an element $g \in \mathrm{Aff}(\mathbb{R}^d)$ is in G if and only if there exists $g_t \in G_t$ for each $t>0$ such that $g_t \to g$ as $t \to 0$.

Proposition 3.8. *The Lie subgroup $G \subset \mathrm{Aff}(\mathbb{R}^d)$ has the following properties.*

1. $\dim G = \dim G_t$ *for all* $t>0$.
2. *Each $g \in G$ preserves the flat degenerate metric determined by $\mathscr{Q}$.*
3. *Each $g \in G$ preserves the fibration determined by the projection $\pi : \mathbb{R}^d = \mathbb{R}^m \oplus \mathbb{R}^{d-m} \to \mathbb{R}^m$. If Q_V denotes the restriction of Q to $\mathbb{R}^m$, a non-degenerate form, then g acts as a Q_V-isometry, denoted $\varpi(g)$, on $\mathbb{R}^m$. The map $\varpi : G \to \mathrm{Isom}(Q_V)$ is a surjective homomorphism satisfying the equivariance property: $\pi(gx) = \varpi(g)\pi(x)$.*

(4) *Let Q_U denote the restriction of Q to the $\mathbb{R}^{d-m}$ factor, a non-degenerate quadratic form. Then Q_U determines, by parallel translation, a flat semi-Riemannian metric on each fiber $\pi^{-1}(y) = \{(y, z) : z \in \mathbb{R}^{d-m}\}$. These semi-Riemannian metrics are also preserved by G, meaning that $g \in G$ takes the metric on $\pi^{-1}(y)$ to the metric on $g\pi^{-1}(y) = \pi^{-1}(\varpi(g)y)$.*
(5) *G is the unique subgroup of $\mathrm{Aff}(\mathbb{R}^d)$ satisfying* (3) *and* (4).

In coordinates respecting the splitting $\mathbb{R}^d = \mathbb{R}^m \oplus \mathbb{R}^{d-m}$, we may describe the group G as the collection of all affine transformations with any translational part and whose linear part has the form

$$\begin{pmatrix} A & 0 \\ B & C \end{pmatrix}$$

where $A \in \mathrm{O}(Q_V), C \in \mathrm{O}(Q_U)$ and B is any $(d-m) \times m$ matrix.

The geometry defined by the group G is much more rigid then the geometry defined by just the degenerate flat metric associated to $\mathscr{Q}$. This more rigid geometry appears naturally in the study of geometric structures transitioning from flat semi-Riemannian geometry of one signature to another, see [5]. While the notion of Delaunay decomposition with respect to $\mathscr{Q}$ is *not* preserved by $\mathrm{Isom}(\mathscr{Q})$, it is preserved by G. As in the non-degenerate case, see Section 3.4 for a stronger result about the invariance of the $\mathscr{Q}$–Delaunay decompositions.

Proposition 3.9. *The set of $\mathscr{Q}$-balls is invariant under G. Hence so is the notion of $\mathscr{Q}$-Delaunay decomposition: if $g \in G$ and $X \subset \mathbb{R}^d$ is in spacelike and generic position with respect to $\mathscr{Q}$, then* $\mathrm{Del}_{\mathscr{Q}}(g(X)) = g(\mathrm{Del}_{\mathscr{Q}}(X))$.

Proof. We use the notations introduced in Lemma 3.7 and Proposition 3.8 above. For any $t > 0$, a simple calculation shows that the map L_t takes Q-balls to Q_{t^2}-balls (bijectively). Hence Lemma 3.7 implies that every $\mathscr{Q}$-ball is a limit of Q_t-balls. The result follows because G_t preserves the set of all Q_t balls so its limit group G must preserve the set of limits of Q_t balls which includes the $\mathscr{Q}$ balls. □

Remark 3.10. There does not seem to be any reasonable 'geometric' notion of Voronoi decomposition with respect to a degenerate quadratic form $\mathscr{Q}$, because the $\mathscr{Q}$-balls do not have a center: one should think of the center as being at infinity. However, we note that by Theorem 3.6, the combinatorics of $\mathrm{Vor}_Q(X_t)$ is constant, so it may be natural to define a combinatorial $\mathscr{Q}$–Voronoi decomposition to agree with the combinatorics of $\mathrm{Vor}_Q(X_t)$.

3.3. **The convex hull construction revisited.** Let us now return to the convex hull construction used in the proof of Theorem A. First, let Q be a non-degenerate quadratic form on $\mathbb{R}^d$ of signature (p, q). Consider the quadratic form $\widehat{Q}$ on $\mathbb{R}^{d+2}$ defined by

$$\widehat{Q}(x_1, \ldots, x_d, x_{d+1}, x_{d+2}) := Q(x_1, \ldots, x_d) - x_{d+1}x_{d+2}.$$

Then $\widehat{Q}$ has signature $(p+1, q+1)$, and the open subset $\mathbb{X} \subset \mathbb{RP}^{d+1}$ consisting of negative lines in $\mathbb{R}^{d+2}$ with respect to $\widehat{Q}$, that is

$$\mathbb{X} = \{x \in \mathbb{R}^{d+2} \setminus \{0\} \mid \widehat{Q}(x) < 0\}/\mathbb{R}^*,$$

is a model for semi-Riemannian geometry of signature $(p+1, q)$ with constant negative curvature. The group $\mathrm{PO}(\widehat{Q})$ of linear transformations of $\mathbb{RP}^{d+1}$ preserving $\widehat{Q}$ is the isometry group of a homogeneous semi-Riemannian metric of signature $(p+1, q)$ and constant negative sectional curvature. Now, consider the affine chart $x_{d+2} = 1$. In this chart, the ideal boundary of $\mathbb{X}$ is described by the equation

$$x_{d+1} = Q(x_1, \ldots, x_d),$$

i.e. the graph in $\mathbb{R}^{d+1}$ of Q. Hence, via the map $x \mapsto (x, Q(x))$, $\mathbb{R}^d$ may be regarded as an open (and dense) set on the ideal boundary $\partial\mathbb{X} \subset \mathbb{RP}^{d+1}$. Indeed, $\partial\mathbb{X}$ is the natural conformal compactification of the flat semi-Riemmanian metric on $\mathbb{R}^d$ defined by Q, and the affine isometries $\mathrm{O}(Q) \ltimes \mathbb{R}^d$ of that flat metric on $\mathbb{R}^d$ are naturally a subgroup of the full conformal group $\widehat{G} = \mathrm{PO}(\widehat{Q})$ of $\partial\mathbb{X}$. Hence in the proof of Theorem A, the convex hull $\mathscr{C}$ of the lifts X' of X to the graph of Q may be regarded as a convex polyhedron in $\mathbb{X}$ whose vertices lie on the ideal boundary $\partial\mathbb{X}$. Such a polyhedron $\mathscr{C}$ is called an *ideal polyedron*.

Next, suppose $\mathbb{R}^d = \mathbb{R}^m \oplus \mathbb{R}^{d-m}$ is a Q–orthogonal decomposition, with $Q = Q_V \oplus Q_U$ and, as in Section 3.2, consider the degenerate quadratic form $\mathscr{Q}$ on $\mathbb{R}^d$ defined by $\mathscr{Q}(y, z) = Q_V(y)$ for all $y \in \mathbb{R}^m, z \in \mathbb{R}^{d-m}$. In Section 3.2, we defined a subgroup G of the group of affine transformations of $\mathbb{R}^d$ which captures the geometry of a quadratic form thought of as having finite part the degenerate quadratic form $\mathscr{Q}$ and infinitesimal part described by Q_U in the degenerate directions $\mathbb{R}^{d-m}$. The group G was defined to be the limit as $t \to 0$ of the isometry groups of the flat metrics defined by quadratic forms $Q_t = Q_V \oplus tQ_U$. We may similarly consider the quadratic forms

$$\widehat{Q}_t(x_1, \dots, x_{d+2}) = Q_V(x_1, \dots, x_m) + tQ_U(x_{m+1}, \dots, x_d) - x_{d+1}x_{d+2}.$$

We denote the limit as $t \to 0$ of the projective orthogonal groups $\mathrm{PO}(\widehat{Q}_t)$ by $\widehat{G}$. Then $\widehat{G}$ acts on the space $\mathbb{X} \subset \mathbb{RP}^{d+1}$ of lines of negative signature with respect to $\widehat{\mathscr{Q}} = \mathscr{Q} - x_{d+1}x_{d+2}$ preserving a degenerate semi-Riemannian metric defined by $\mathscr{Q}$ and also preserving a quadratic form naturally isomorphic to Q_U on the degenerate subspace, a copy of $\mathbb{R}^{d-m}$, of each tangent space. Note that G is naturally the subgroup of $\widehat{G}$ that preserves the affine chart $x_{d+2} = 1$. As in the non-degenerate case above, the subset of the ideal boundary $\partial\mathbb{X}$ that lies in the affine chart $x_{d+2} = 1$ naturally identifies with $\mathbb{R}^d$ via the map $x \mapsto (x, \mathscr{Q}(x))$, and the full ideal boundary $\partial\mathbb{X}$ is thought of as a conformal compactification of the degenerate geometry of $\mathbb{R}^d$ described by G. Again, the points of X in Theorem A are thought of as points on $\partial\mathbb{X}$ and the convex hull $\mathscr{C}$ constructed in the proof of the theorem is an ideal polyhedron in $\mathbb{X}$. The space $\mathbb{X}$ and its symmetry group $\widehat{G}$ describe constant curvature semi-Riemannian geometry which is infinitesimal in some directions. See [5] for more on constant curvature semi-Riemannian geometries and their limits.

In dimension $d = 2$, consider the Euclidean quadratic form $Q(x_1, x_2) = x_1^2 + x_2^2$. Then $\widehat{Q}(x_1, x_2, x_3, x_4) = x_1^2 + x_2^2 - x_3x_4$ defines a copy $\mathbb{X}$ of the three-dimensional hyperbolic space $\mathbb{H}^3$. The Euclidean plane $\mathbb{R}^2$ naturally identifies with the subset of the ideal boundary $\partial\mathbb{X}$ lying in the affine patch $x_4 = 1$. A set of points X in spacelike and general position with respect to Q determines a convex ideal polyhedron $\mathscr{C}$ in $\mathbb{H}^3$ and conversely.

Similarly, consider the standard Lorentzian quadratic form $Q(x_1, x_2) = x_1^2 - x_2^2$. Then $\widehat{Q}(x_1, x_2, x_3, x_4) = x_1^2 - x_2^2 - x_3x_4$ defines a copy $\mathbb{X}$ of the $2+1$ dimensional *anti de Sitter space* $\mathbb{A}\mathrm{d}\mathbb{S}^3$, a model for constant negative curvature Lorentzian geometry in dimension $2 + 1$. The Minkowski space $\mathbb{R}^{1,1}$ naturally identifies with the subset of the ideal boundary $\partial\mathbb{X}$ lying in the affine patch $x_4 = 1$. Again, a set of points X in spacelike and general position with respect to Q determines a convex ideal polyhedron $\mathscr{C}$ in $\mathbb{A}\mathrm{d}\mathbb{S}^3$ and conversely.

Next, define the degenerate quadratic form $\mathscr{Q}$ from Q being either the Euclidean quadratic form or the Lorentzian quadratic form above using the coordinate splitting $\mathbb{R}^2 = \mathbb{R}^1 \oplus \mathbb{R}^1$ and let $\widehat{G}$ and $\mathbb{X}$ be defined as above. Then the degenerate plane $\mathbb{R}^{1,0,1}$ naturally identifies with the subset of the ideal boundary $\partial\mathbb{X}$ lying in the affine chart $x_4 = 1$. Here $\mathbb{X}$ and its symmetry group $\widehat{G}$ gives what is known as three-dimensional *half-pipe geometry* $\mathbb{HP}^3$, defined in [6]. It may be regarded

as the geometry of an infinitesimal neighborhood of a hyperbolic 2-plane in either the three-dimensional hyperbolic space $\mathbb{H}^3$ or the three-dimensional anti de Sitter space $\mathbb{A}\mathrm{d}\mathbb{S}^3$. Again, a set of points X in spacelike and general position with respect to $\mathscr{Q}$ determines a convex ideal polyhedron $\mathscr{C}$ in $\mathbb{HP}^3$ and conversely.

In Section 4.3, we will relate the geometry of the Delaunay triangulation of X with the geometry of the corresponding ideal polyhedron $\mathscr{C}$ in each of the three geometric contexts above. We then characterize Delaunay triangulations in $\mathbb{R}^{1,1}$ and $\mathbb{R}^{1,0,1}$ by applying a recent characterization of ideal polyhedra that we obtained in [7]. The following theorem extends to $\mathbb{A}\mathrm{d}\mathbb{S}^3$ and $\mathbb{HP}^3$ a famous theorem of Rivin [15] about ideal polyhedra in hyperbolic three-space.

Theorem 3.11 ([7], Theorem 1.3 and 1.7)**.** *Let $\mathbb{X}$ be $\mathbb{A}\mathrm{d}\mathbb{S}^3$ or $\mathbb{HP}^3$. Let Γ' be a triangulated graph on the sphere and let $w : E(\Gamma') \to \mathbb{R} \setminus \{0\}$ be a weight function on the edges $E(\Gamma')$ of Γ'. Then there exists a convex ideal polyhedron $\mathscr{C}$ in $\mathbb{X}$ and an isomorphism of Γ' with the one-skeleton of $\mathscr{C}$ which takes the weights w to the dihedral angles of $\mathscr{C}$ if and only if the following conditions hold:*

(i) For each vertex v of Γ', the vertex sum $\sum_{e\sim v} w(e) = 0$.
(ii) The edges e for which $w(e) < 0$ form a Hamiltonian cycle in Γ'.
(iii) If c is a Jordan curve on the sphere transverse to Γ' which crosses exactly two edges of negative weight, then

$$\sum_{e\in E_c} w(e) \geq 0$$

where E_c are the edges of Γ' crossed by c. Further, equality occurs if and only if c encircles a single vertex.

3.4. **Möbius invariance of Q–Delaunay decompositions.** Propositions 3.1 and 3.9 state that Q–Delaunay decompositions are invariant under the action of the isometry group of the flat geometry associated to Q. In this section, we search for a larger collection of transformations leaving the Q–Delaunay decomposition invariant. To this end, recall that the construction of the Q–Delaunay decomposition for X in the proof of Theorem A, interpreted according to the observations of Section 3.3, involves forming the convex ideal polyhedron $\mathscr{C}$ in the negatively curved geometry $\mathbb{X}$ associated to $\widehat{Q}$ whose vertices are the points X, thought of as lying in the ideal boundary $\partial\mathbb{X}$ using the standard chart $\mathbb{R}^d \hookrightarrow \partial\mathbb{X}$. The convex hull is a construction invariant under the full group of Q-Möbius transformations $\widehat{G} = \mathrm{O}(\widehat{Q})$. However, the second half of the construction of $\mathrm{Del}_Q(X)$ involves projecting the *bottom faces* $\partial\mathscr{C}_-$ back down to $\mathbb{R}^d$ and the notion of bottom and top faces of $\mathscr{C}$ depends on a choice of a point $\infty \in \partial\mathbb{X}$ to project from. So, the Q–Delaunay decomposition of X should be invariant under any Möbius transformation $f \in \widehat{G}$ for which the bottom faces of $f(\mathscr{C})$ relative to ∞ are precisely the image of the bottom faces of $\mathscr{C}$ relative to ∞. Let us now make this more precise.

First, we note that strictly speaking, even for small $f \in \widehat{G}$, $f(\mathrm{Del}_Q(X))$ is rarely a *linear* cellulation; indeed, the image of an affine plane under f is typically some Q-sphere (of the same dimension). However, as long as the image of a cell of $\mathrm{Del}_Q(X)$ is contained (and compact) in $\mathbb{R}^d$, we may *isotop* it relative to its vertices so that it becomes a linear cell (i.e. pull it tight). Note that even if the image of each cell of $\mathrm{Del}_Q(X)$ lies in $\mathbb{R}^d$, the straightening of $f(\mathrm{Del}_Q(X))$ may fail to be a cellulation of the convex hull of $f(X)$ (and even if it is a cellulation of the convex hull, it might not be $\mathrm{Del}_Q(f(X))$). Henceforth we will write $f(\mathrm{Del}_Q(X)) \sim \mathrm{Del}_Q(f(X))$ to mean that $\mathrm{Del}_Q(f(X))$ is the straightening of $f(\mathrm{Del}_Q(X))$.

As a warm-up to the general case, we first discuss the Euclidean case, where the precise statement of the Möbius invariance (below) is presumably well-known

to specialists. Given X, let $\mathcal{C}(X)$ denote the finite collection of all full and empty spheres (see Definition 2.3 and Remark 2.4).

Lemma 3.12. *Let X be a finite set of points in $\mathbb{R}^d$, let Q be the standard Euclidean quadratic form and let $f \in \widehat{G} = \mathrm{PO}(d,1)$. Then $f(\mathrm{Del}_Q(X))$ is isotopic to $\mathrm{Del}_Q(f(X))$ if and only if $f^{-1}(\infty)$ lies outside each sphere of the collection $\mathcal{C}(X)$ of empty and full spheres associated to X.*

Remark 3.13. Note that the conclusion of Lemma 3.12 may fail if the condition is only required to hold for the empty circles. There are easy examples of this even with $|X| = 4$.

Outline of the proof. Let $f \in \widehat{G}$ and assume that $f(X)$ lies in $\mathbb{R}^d$ (no point of X is sent to ∞). Consider a sphere S in $\mathbb{R}^d$. The condition that S be empty of points of X means that in the conformal compactification $\partial\mathbb{X} = \mathbb{R}^d \cup \{\infty\}$, ∞ lies on one side of S while X lies on the other (allowing for some points of X to be on S). Hence if S is an empty sphere, then $f(S)$ is empty of points of $f(X)$ if and only if $f^{-1}(\infty)$ lies in the component of the complement of S containing ∞, i.e. $f^{-1}(\infty)$ is outside of S. Hence, if $f^{-1}(\infty)$ lies outside all of the empty spheres associated to X, then the cells of $f(\mathrm{Del}(X))$ straighten to cells of $\mathrm{Del}(f(X))$. However, $\mathrm{Del}(f(X))$ could have additional top-dimensional cells, if there are additional empty spheres. Since the top-dimensional cells of $\mathrm{Del}(f(X))$ are projections of faces of the convex hull $f(\mathscr{C})$ of the points $f(X) \subset \partial\mathbb{X}$ (using here that the convex hull construction in $\mathbb{X}$ is invariant under $\widehat{G}$), we need only check now that the images of the full spheres associated to X are not empty for $f(X)$. This is precisely the condition that $f^{-1}(\infty)$ lies outside of every full sphere associated to X. □

The condition in Lemma 3.12 does not generalize well to the case of non-positive definite quadratic forms, since in general a Q–sphere in $\partial\mathbb{X}$ does not separate $\partial\mathbb{X}$ making it non-sensical to ask for a point of $\partial\mathbb{X}$ to lie inside or outside of a Q-sphere. For the sake of generalizing, let us reformulate Lemma 3.12 as follows.

Lemma 3.14. *Let X be a finite set of points in $\mathbb{R}^d$, let Q be the standard Euclidean quadratic form, and let $f \in \widehat{G} = \mathrm{PO}(d,1)$. Then $f(\mathrm{Del}_Q(X))$ is isotopic to $\mathrm{Del}_Q(f(X))$ if and only if $f^{-1}(\infty)$ is in the same connected component as ∞ in the complement of the union of the spheres in $\mathcal{C}(X)$ in the conformal compactification $\partial\mathbb{X}$ of $(\mathbb{R}^d, Q)$.*

In the general setting, the condition above needs to be slightly expanded as follows. Recall from Section 3.3 that, for a general quadratic form Q, the conformal compactification $\partial\mathbb{X}$ contains many points at infinity, indeed $\partial\mathbb{X} = \mathbb{R}^d \cup \mathscr{L}(\infty)$ is the disjoint union of $\mathbb{R}^d$ with the light cone of a point ∞ at infinity. Here the light cone $\mathscr{L}(x)$ of a point $x \in \partial\mathbb{X}$ is the union of all $y \in \partial\mathbb{X}$ such that $\langle x, y\rangle_{\hat{Q}} = 0$. If $x \in \mathbb{R}^d$, then $\mathscr{L}(x) \cap \mathbb{R}^d$ is precisely the points of $\mathbb{R}^d$ of null displacement from x in the Q norm (this only looks like a standard cone if Q is Lorentzian). The stereographic projection used in the proof of Theorem A then conformally identifies the complement of $\mathscr{L}(\infty)$ in $\partial\mathbb{X}$ with $(\mathbb{R}^d, Q)$.

Theorem 3.15. *Let Q be any quadratic form on $\mathbb{R}^d$ and let X be a finite set of points in $\mathbb{R}^d$ in space-like position with respect to Q. Let $f \in \widehat{G}$. Then $f(\mathrm{Del}_Q(X))$ is isotopic to $\mathrm{Del}_Q(f(X))$ if $f^{-1}(\infty)$ lies in the same connected component as ∞ in the complement in the conformal compactification $\partial\mathbb{X}$ of $(\mathbb{R}^d, Q)$ of the union of*

(1) the collection $\mathcal{C}$ of all the empty and full Q-spheres,
(2) the light cones $\mathscr{L}(x)$ of the points of X.

The proof of Theorem 3.15 uses a simple statement in projective geometry. We use the notation of Section 3.3.

Lemma 3.16. *Let S be a Q-sphere in $\partial\mathbb{X}$ realized as the (transverse) intersection of a hyperplane $P \subset \mathbb{RP}^{d+1}$ with $\partial\mathbb{X}$. Let $P_0 = P \cap \mathbb{X}$. Then a point $x \in \mathbb{R}^d \setminus S$ lies inside the Q-sphere S if and only if the line in $\mathbb{RP}^{d+1}$ passing through x and ∞ crosses the hyperplane P inside of P_0.*

Proof. Let S be defined by the equation:

$$Q(x) = \varphi(x) + D$$

for $x \in \mathbb{R}^d$, where $\varphi : \mathbb{R}^d \to \mathbb{R}$ is a linear functional and $D \in \mathbb{R}$ is a constant. Recall the explicit embedding $x \mapsto [x : Q(x) : 1] \in \mathbb{RP}^{d+1}$. Then P is the set of $[y : a : b] \in \mathbb{RP}^{d+1}$ such that $\varphi(y) = a - Db$, where $y \in \mathbb{R}^d$ and $a, b \in \mathbb{R}$. The point $\infty = [0_d : 1 : 0]$ (where 0_d is the zero vector in $\mathbb{R}^d$) and let $x \in \mathbb{R}^d \setminus S$. Then the line ℓ in $\mathbb{RP}^{d+1}$ passing through ∞ and $x$is given by $[tx : tQ(x) + s : t]$. Then ℓ passes through P at the point p where $s = t(\varphi(x) - Q(x) + D)$. Let us evaluate (the sign of) $\widehat{Q}$ at this point:

$$\begin{aligned}\widehat{Q}(tx, tQ(x) + s, t) &= \widehat{Q}(tx, tQ(x) + t(\varphi(x) - Q(x) + D), t)\\ &= \widehat{Q}(tx, t\varphi(x) + tD, t)\\ &= Q(tx) - (t\varphi(x) + tD)t = t^2(Q(x) - (\varphi(x) + D)).\end{aligned}$$

Hence $p \in \mathbb{X}$ if and only if $Q(x) < \varphi(x) + D$ if and only if x is inside the Q-sphere S. □

Proof of Theorem 3.15. Let $f_t \in \widehat{G}$ be a family of Möbius transformations with $f_0 = I$ such that for all t, $f_t^{-1}(\infty)$ does not lie on the light cone of any point of X not on any empty or full Q-sphere in $\mathcal{C}$. We must first show that $f(X)$ is spacelike position, so that $\mathrm{Del}_Q(f(X))$ is defined. Consider the spacelike segment σ_{ij} connecting x_i to x_j in $\mathbb{R}^d$. Then, since f_t is a conformal deformation, $f_t(\sigma_{ij})$ is a spacelike segment in the conformal metric on $\partial\mathbb{X}$ for all t. The only way $f_t(x_i)$ and $f_t(x_j)$ could fail to have spacelike relative position is if $f_t(\sigma_{ij})$ is not contained in $\mathbb{R}^d$, equivalently if $f_t(\sigma_{ij})$ crosses $\mathscr{L}(\infty)$, equivalently if $f_t^{-1}\mathscr{L}(\infty) = \mathscr{L}(f_t^{-1}(\infty))$ crosses σ_{ij}. If this happens, then for some (possibly smaller) value of t, $\mathscr{L}(f_t^{-1}(\infty))$ contains the point x_i or x_j, equivalently $f_t^{-1}\infty$ lies in $\mathscr{L}(x_i)$ or $\mathscr{L}(x_j)$. This does not happen by assumption.

Now, to show that $f_t(\mathrm{Del}_Q(X)) \sim \mathrm{Del}_Q(f_t(X))$, we argue as in the Euclidean case. Let Δ be a cell of $\mathrm{Del}_Q(f_t(X))$. Then Δ is the projection to $\mathbb{R}^d$ of a face of the convex hull in $\mathbb{X}$ of the points $f_t(X) \subset \partial\mathbb{X}$. This convex hull is precisely $f_t(\mathscr{C})$, hence the vertices of Δ are the images under f_t of a subset of X which span a top or bottom face of $\mathscr{C}$, and therefore define one of the empty or full spheres with respect to X. So it remains to show that for each empty (respectively full) sphere S in $\mathcal{C}$, $f_t(S)$ remains empty (respectively full) with respect to $f_t(X)$.

Consider an empty or full Q-sphere S in $\mathcal{C}(X)$. We now show that $f_t(S)$ is an empty (resp. full) sphere for $f_t(X)$ for all t. Let $P \subset \mathbb{RP}^{d+1}$ be the hyperplane as in Lemma 3.16 such that $P \cap \partial\mathbb{X} = S$. Let $x_i \in X$ be a point not on S. By the Lemma, the point $f_t(x_i)$ lies strictly inside the Q-sphere $f_t(S)$ if and only if the line ℓ_t passing through $f_t(x_i)$ and ∞ intersects $f_t(P)$ in the interior of $\mathbb{X}$. If $f_{t'}(x_i)$ is at some time t' on the wrong side of $f_{t'}(S)$, then there is a time t for which $\ell_t \cap f_t(P) \in \partial\mathbb{X} \cap f_t(P) = f_t(S)$, or equivalently the line $f_t^{-1}\ell_t$ in $\mathbb{RP}^{d+1}$ passing through x_i and $f_t^{-1}(\infty)$ intersects S. Hence, since $x_i \notin S$, then either $f_t^{-1}(\infty) \in S$, which we assume does not happen, or the entire line $f_t^{-1}\ell_t$ lies in $\partial\mathbb{X}$ which implies that $f_t^{-1}(\infty) \in \mathscr{L}(x_i)$, which we also assume does not happen. Thus for any t, $f_t(x_i)$ lies strictly inside $f_t(S)$ if and only if x_i lies strictly inside S. □

4. Higher signature Delaunay decompositions in dimension $d = 2$

In dimension $d = 2$, there are two interesting geometric settings to consider beyond the classical case of the Euclidean plane (i.e. $Q = Q_{2,0}$ a positive definite form). The first is the *Minkowski plane*, $\mathbb{R}^{1,1}$, equipped with quadratic form $Q_{1,1}(x_1, x_2) = x_1^2 - x_2^2$. See Section 3.1. The second is the degenerate geometry $\mathbb{R}^{1,0,1}$ defined by the degenerate quadratic form $Q_{1,0,1}(x_1, x_2) = x_1^2$ as in Sections 3.2 and Section 3.3.

There are two collections of angles associated to a Q–Delaunay triangulation $\mathrm{Del}_Q(X)$, the collection of *interior angles* (Definition 4.4) of the triangles and the collection of *edge angles*, which are angles formed by the Q-circles circumscribing adjacent triangles (Definition 4.5), assigned to the edges of $\mathrm{Del}_Q(X)$. We investigate the behavior of the interior angles and edge angles generalizing several important results from the Euclidean setting.

4.1. Interior angles and edge angles of higher signature Delaunay triangulations.

First, we define a notion of angle in each of the two geometries of interest.

4.1.1. *Angles in the Minkowski plane* $\mathbb{R}^{1,1}$. In the Minkowski plane $\mathbb{R}^{1,1}$, there are two connected components of spacelike directions. The angle formed by two spacelike unit vectors v, w in the same component is defined to be the positive number $\varphi > 0$ satisfying the equation

$$\cosh \varphi = \langle v, w \rangle_{Q_{1,1}}.$$

The angle formed by two spacelike unit vectors v, w in opposite components is defined to be the negative number $\varphi < 0$ satisfying

$$\cosh \varphi = -\langle v, w \rangle_{Q_{1,1}}.$$

4.1.2. *Angles in the degenerate plane* $\mathbb{R}^{1,0,1}$. Following the notation and ideas of Section 3.2, let $Q_{1,0,1} = Q_V$ and $Q_{2,0} = Q_V + Q_U$ and $Q_{1,1} = Q_V - Q_U$, where $Q_V(x_1, x_2) = x_1^2$ and $Q_U(x_1, x_2) = x_2^2$. Here we focus on the case of $Q_{1,0,1}$. We say that a non-zero vector v is *non-degenerate* if $Q_{1,0,1}(v, v) > 0$.

Let $v = (v_1, v_2), w = (w_1, w_2)$ be two unit vectors with respect to $Q_{1,0,1}$, i.e. $v_1^2 = w_1^2 = 1$. Recall the map $L_t : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $L_t(x_1, x_2) = (x_1, \frac{x_2}{t})$ and consider any two paths $v_t = (a_t, b_t), w_t = (c_t, d_t)$ so that $L_t v_t \to v$ and $L_t w_t \to w$ as $t \to 0$, i.e. $a_t \to v_1, b_t \to 0$, $c_t \to w_1, d_t \to 0$ and $\dot{b} := \frac{\mathrm{d}}{\mathrm{d}t}\big|_{t=0} b_t = v_2$ and $\dot{d} := \frac{\mathrm{d}}{\mathrm{d}t}\big|_{t=0} d_t = w_2$. Then, we define the angle $\dot{\theta}_{vw}$ between v and w in $\mathbb{R}^{1,0,1}$ to be the derivative

$$\dot{\theta}_{vw} := \frac{\mathrm{d}}{\mathrm{d}t}\big|_{t=0} \theta_{v_t w_t},$$

where $\theta_{v_t w_t}$ denotes the $\mathbb{R}^{2,0}$ or $\mathbb{R}^{1,1}$ angle between v_t and w_t. It is easy to check that this definition does not depend on the paths of vectors v_t, w_t chosen, nor does it depend on whether $\theta_{v_t w_t}$ is measured in $\mathbb{R}^{2,0}$ or $\mathbb{R}^{1,1}$ because in all cases

$$\dot{\theta}_{vw} = \begin{cases} |v_2 - w_2| & \text{if } v_1 w_1 = 1 \\ -|v_2 - w_2| & \text{if } v_1 w_1 = -1 \end{cases}. \tag{6}$$

Proposition 4.1. *The angle $\dot{\theta}_{vw}$ between two directions v, w is invariant under the symmetry group G of $\mathbb{R}^{1,0,1}$ defined in Section 3.2.1.*

Proof. Let $v, w \in \mathbb{R}^{1,0,1}$, let $g \in G$, and let $v' = gv, w' = gw'$. Let $(v_t)_{t\in(0,1)}$ and $(w_t)_{t\in(0,1)}$ be chosen such that $L_t v_t \to v$ and $L_t w_t \to w$ as $t \to 0$, as above. Let

$(g_t)_{t\in(0,1)}$, $g_t \in G_{t^2}$, be such that $g_t \to g$, where G_{t^2} was defined in Section 3.2.1. Then $g_t L_t v_t \to v'$, $g_t L_t w_t \to w'$. Hence since $\mathcal{L} = L_{1/t} g_t L_t$ preserves Q, we have:

$$\theta_{v'w'} = \frac{d}{dt}_{|t=0} \theta_{\mathcal{L}v_t, \mathcal{L}w_t} = \frac{d}{dt}_{|t=0} \theta_{v_t, w_t} = \theta_{vw} \ .$$

□

Proposition 4.2. *Consider a non-degenerate triangle $Q_{1,0,1}$. Then the sum of the $Q_{1,0,1}$ interior angles of the triangle is zero.*

Proof. Take the derivative of the analogous formula in Euclidean or Minkowski geometry. □

Proposition 4.3. *Let p_1, p_2 be two points with non-degenerate displacement and let p_3, p_4 be points in $\mathbb{R}^{1,0,1}$ lying in the same connected component of the set of points which are non-degenerate with respect to both p_1 and p_2. Suppose that p_4 is in the interior of the triangle $\Delta p_1 p_2 p_3$. Then the $\mathbb{R}^{1,0,1}$ angles $\measuredangle p_1 p_3 p_2, \measuredangle p_1 p_4 p_2$ satisfy that*

$$\measuredangle p_1 p_3 p_2 < \measuredangle p_1 p_4 p_2.$$

Proof. This follows from the definition of the $\mathbb{R}^{1,0,1}$ angles in terms of limits of angles between triples of points in $\mathbb{R}^2$ (or $\mathbb{R}^{1,1}$), and from the corresponding inequality in $\mathbb{R}^2$. □

4.2. **Interior angles and edge angles.** Let $Q = Q_{1,1}$ or $Q_{1,0,1}$. There are two collections of angles naturally assigned to a Q-Delaunay triangulation $\mathrm{Del}_{Q_{1,1}}(X)$. Each triangle has three *interior angles*, two positive and one negative, which sum to zero.

Definition 4.4. Let $Q = Q_{1,1}$ or $Q = Q_{1,0,1}$ and let X be a finite set in $\mathbb{R}^2$ in spacelike and generic position with respect to Q. The *interior angles* of $\mathrm{Del}_Q(X)$ is the collection of interior angles of all triangles of $\mathrm{Del}_Q(X)$, listed in increasing order.

Each edge e of $\mathrm{Del}_{Q_{1,1}}(X)$ is assigned an *edge angle* defined in terms of the intersection between the empty Q-circles circumscribing the triangles adjacent to that edge. To make a precise definition, consider two regions R_1, R_2 in the plane, each with smooth boundary $S_1 = \partial R_1$, $S_2 = \partial R_2$. The angle formed by R_1 and R_2 at a point p in the intersection $S_1 \cap S_2$ of their boundaries is defined to be the angle φ formed by the two unit vectors v_1, v_2 tangent to S_1 and S_2 at p and in the direction of traversal placing R_1 and R_2 on the left with respect to some (any) fixed ambient orientation. In the case that $R_1 = B_1$, $R_2 = B_2$ are Q–balls, the corresponding Q–circles S_1 and S_2 will typically intersect at two points. It is easy to check that the angle formed by B_1 and B_2 at both points is the same, making the notion of intersection angle between Q–circles well defined.

Definition 4.5. Let X be a finite set of points in $\mathbb{R}^2$ in spacelike and general position with respect to a quadratic form Q and let $\mathcal{T}$ be any triangulation of the convex hull of X. Consider an edge e of $\mathcal{T}$. If e is an interior edge, then it is adjacent to two triangles t_1, t_2. In this case, the *edge angle* at e of $\mathcal{T}$ is the angle φ formed by the Q balls B_1 and B_2 circumscribing t_1 and t_2 respectively. If e is an exterior edge, then e is adjacent to one triangle t_1. In this case, the *edge angle* at e is the angle between the ball B_1 circumscribing t_1 and the half-plane whose intersection with the convex hull of X is precisely e. Typically we will denote by $\Phi : E(\mathcal{T}) \to \mathbb{R}$ the function which assigns to each edge e in the set of edges $E(\mathcal{T})$ the dihedral angle $\Phi(e)$ of $\mathcal{T}$ at e.

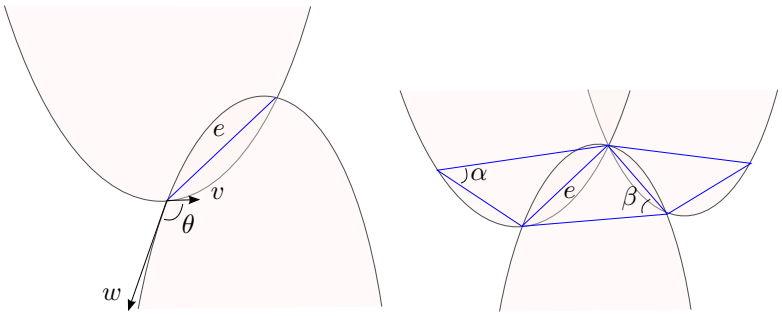

FIGURE 4. The edge angle θ associated to e is the angle between the unit vectors v and w tangent to the two empty Q-circles circumscribing the triangles adjacent to e. Alternatively, θ is the sum of the interior angles α, β opposite to e.

As for Euclidean Delaunay decompositions, there is a simple relation between the interior angles and the edge angles:

Proposition 4.6. *Let e be an edge of $\mathrm{Del}_Q(X)$. Then the edge angle at e is the sum of the interior angles opposite to e in the triangles adjacent to e. If e is an interior edge, then there are two such triangles, whereas if e is an exterior edge, than there is one such triangle.*

In the case $Q = Q_{1,1}$, the proof of this property follows from an elementary geometric remark: given three points a, b, c at time-distance $r > 0$ from 0 in $\mathbb{R}^{1,1}$, then the Minkowski angles satisfy $\measuredangle(bac) = \measuredangle(b0c)/2$. In the case $Q = Q_{1,0,1}$, the proof proceeds by taking limits of the Minkowski or Euclidean case as usual.

4.3. **Prescribing the edge angles.** The edge angles of any triangulation $\mathcal{T}$ as in Definition 4.5 determines a weighted planar triangulated graph. In the case that the quadratic form Q is not positive definite, the following theorem characterizes exactly which weighted graphs occur as the edge angles of a Q–Delaunay triangulation.

Theorem 4.7. *Let Γ be a triangulated graph in the plane and let $w : E(\Gamma) \to \mathbb{R}\backslash\{0\}$ be a weight function on the set $E(\Gamma)$ of edges of Γ. Let Q be either $Q_{1,1}$ or $Q_{1,0,1}$. Then there exists a finite set X in $\mathbb{R}^2$ in spacelike and general position with respect to Q and an isomorphism of Γ with the one-skeleton of the Delaunay triangulation $\mathrm{Del}_Q(X)$ which takes the weights w to the edge angles of $\mathrm{Del}_Q(X)$ if and only if the following conditions hold:*

1. *For each interior vertex v of Γ, the vertex sum $\sum_{e\sim v} w(e) = 0$.*
2. *The edges e for which $w(e) < 0$ form a Hamiltonian path whose endpoints v_1, v_2 are exterior vertices.*
3. *The vertex sum $\sum_{e\sim v} w(e)$ associated to an exterior vertex v is positive if $v = v_1$ or $v = v_2$ and is negative otherwise.*
4. *If c is a Jordan curve in $\mathbb{R}^2$ transverse to Γ which either crosses two edges of negative weight but does not enclose v_1 or v_2, encloses v_1 and v_2 but does not cross any edge of negative weight, or crosses exactly one edge of negative weight and encloses exactly one of v_1 and v_2, then*

$$\sum_{e\in E_c} w(e) - \sum_{v\in V_c}\sum_{e\sim v} w(e) \geq 0$$

where E_c are the edges of Γ crossed by c and V_c are the exterior vertices of Γ enclosed inside of c. Equality occurs if and only if c encircles a single vertex or c encircles all vertices.

Proof. First, consider a set X of finitely many points in spacelike and generic position with respect to Q. As in the proof of Theorem A, we lift X to a point set X' lying on the graph of Q in $\mathbb{R}^3$. The Delaunay triangulation $\mathrm{Del}_Q(X)$ is seen on the bottom of the convex hull $\mathscr{C}$ of X'. As described in Section 3.3, the graph of Q in $\mathbb{R}^3$ is naturally an open dense subset of the ideal boundary $\partial\mathbb{X}$ of either $\mathbb{X} = \mathbb{A}\mathrm{d}\mathbb{S}^3 \subset \mathbb{RP}^3$ if $Q = Q_{1,1}$ or $\mathbb{X} = \mathbb{HP}^3 \subset \mathbb{RP}^3$ if $Q = Q_{1,0,1}$.

Let us now add an extra vertex $v_\infty = [0 : 0 : 1 : 0]$ and denote the convex hull of X' and v_∞ by $\mathscr{C}_2$. Then the faces of $\mathscr{C}_2$ consist of the bottom faces of $\mathscr{C}$ (corresponding to the triangles of $\mathrm{Del}_Q(X)$) as well as infinite vertical walls, as viewed in $\mathbb{R}^3 \subset \mathbb{RP}^3$ along the edges corresponding to exterior edges of $\mathrm{Del}_Q(X)$. It is straightforward to show that the edge angle at an edge e of $\mathrm{Del}_Q(X)$ (Definition 4.5) is exactly the same as the dihedral angle of the corresponding edge of $\mathscr{C}_2$ as measured in $\mathbb{X}$. Indeed, the conformal semi-Riemannian structure on $\partial\mathbb{X}$ is compatible with the semi-Riemannian structure on $\mathbb{X}$. Further, since the sum of the dihedral angles around an ideal vertex of $\mathscr{C}_2$ are equal to zero (see [7]), the dihedral angle along one of the new vertical edges e of $\mathscr{C}_2$ is exactly minus the sum of the edge angles of the edges of $\mathrm{Del}_Q(X)$ incident to the vertex v to which e projects.

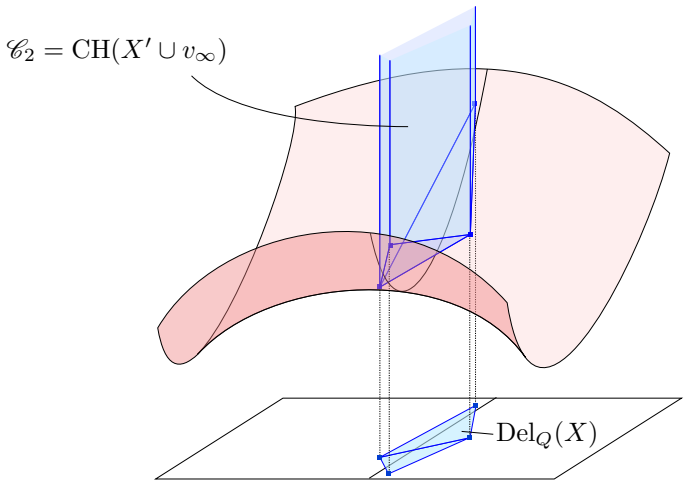

FIGURE 5. I.

Hence, Theorem 4.7 follows directly from Theorem 3.11: Let Γ' denote the triangulated graph on the sphere obtained from Γ by first adding a single vertex v_∞ at infinity and then connecting all exterior vertices of Γ to v_∞ with an edge. There is a one-to-one correspondence between weight functions on Γ satisfying the conditions of Theorem 4.7 and weight functions on Γ' satisfying the conditions of Theorem 3.11. The weights on the new edges of Γ' are determined by Condition (i). Condition (i) for the new vertex v_∞ corresponds to the case that c encircles the entire polygon in Condition (4). Condition (ii) corresponds to Conditions (2) and (3). Finally, in Condition (iii), that c crosses a new edge of Γ' is equivalent to c encircling the corresponding exterior vertex of Γ. □

4.4. **Interior angle optimization.** Consider a triangulation $\mathcal{T}$ of a finite point set X in spacelike and general position with respect to Q, where $Q = Q_{1,1}, Q_{1,0,1}$, or $Q = Q_{2,0}$. The *angle sequence* of $\mathcal{T}$ is the ordered list, sorted from smallest to largest with repetition, of the interior angles of all the triangles in $\mathcal{T}$. Given two triangulations $\mathcal{T}_1, \mathcal{T}_2$, we say that $\mathcal{T}_1$ is *fatter* than $\mathcal{T}_2$ if the angle sequence of $\mathcal{T}_1$ is greater than the angle sequence of $\mathcal{T}_2$ with respect to the lexicographic ordering,

i.e. if the first angle for which the angles sequences of $\mathcal{T}_1$ and $\mathcal{T}_2$ disagree is larger for $\mathcal{T}_1$.

Theorem 4.8. *Let X be a finite set of points in $\mathbb{R}^2$ in spacelike and general position with respect to Q, where $Q = Q_{1,1}, Q_{1,0,1}$, or $Q_{2,0}$. Then the Q–Delaunay triangulation $\mathrm{Del}_Q(X)$ is the unique fattest triangulation of X: the angle sequence of $\mathrm{Del}_Q(X)$ is larger, with respect to the lexicographic ordering, than the angle sequence of any other triangulation of X.*

Of course, this theorem is well-known in the case $Q = Q_{2,0}$ of Euclidean geometry. As in that setting, to prove the theorem, we must first consider the simple case of a quadrilateral. To analyze that case, we need a generalization of the classical Thales' Theorem. The proof in this general setting is essentially the same.

Proposition 4.9 (Thales' Theorem)**.** *Let p_1, p_2, p_3 be three distinct points in spacelike position lying on a Q–circle S and let p_4 and p_5 be points on the same side of the line $\overline{p_1p_2}$ as p_3 and which lie inside and outside of S respectively. Further assume that p_3, p_4, p_5 all lie in the same connected component of the set of points that are in spacelike position with respect to p_1 and p_2. Then $\measuredangle p_1p_4p_2 > \measuredangle p_1p_3p_2 > \measuredangle p_1p_5p_2$.*

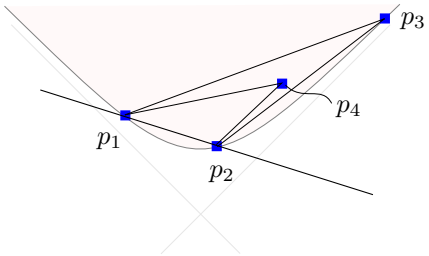

FIGURE 6. In the proof of Proposition 4.9, the angle $\measuredangle p_1p_3p_2$ is locally constant as p_3 is varied, so WLOG we may assume $\Delta p_1p_3p_2$ contains $\Delta p_1p_4p_2$.

Proof. In the case that Q is non-degenerate (so $Q = Q_{2,0}$ or $Q = Q_{1,1}$), the Q–circle has a center O. Observe that the triangles ΔOp_1p_3 and ΔOp_2p_3 are isosceles, from which it easily follows that the angle $\measuredangle p_1p_3p_2$ is locally constant as p_3 is varied along S. So without loss in generality we may assume that $\Delta p_1p_2p_4 \subset \Delta p_1p_2p_3 \subset \Delta p_1p_2p_5$ and the result follows.

In the case $Q = Q_{1,0,1}$, the argument is almost the same. However, in this case a Q–circle S does not have a center (the center is at ∞). Nonetheless, the claim that the angle $\measuredangle p_1p_3p_2$ is locally constant as p_3 is varied along S still holds by taking the derivatives of the same fact for $Q = Q_{2,0}$ as rescaled $Q_{2,0}$-circles converge to S as in Lemma 3.7. So again without loss in generality we may assume that $\Delta p_1p_2p_4 \subset \Delta p_1p_2p_3 \subset \Delta p_1p_2p_5$, and the result follows using the strict inequality of Proposition 4.3. □

Lemma 4.10 (Angle optimization for quadrilaterals)**.** *Consider a convex quadrilateral R whose vertices X are in spacelike general position with respect to Q. Then, of the two triangulations of X, the Q–Delaunay triangulation $\mathrm{Del}_Q(X)$ is the fatter triangulation.*

Proof. The proof is nearly identical to the proof in the classical Euclidean case $Q = Q_{2,0}$. Let p_1, p_2, p_3, p_4 denote the points of X, cyclically ordered around the boundary of R. Without loss in generality, the minimum (which a priori might not

be unique) of the twelve angles formed by any three of the four points is $\measuredangle p_1p_2p_4$. Note that if $Q = Q_{1,1}$ or $Q = Q_{1,0,1}$, then $\measuredangle p_1p_2p_4$ will be negative. In that case, p_3 and p_4 must lie in the same connected component of the set of points in spacelike position with respect to p_1 and p_2. Since p_3 does not lie on the Q–circle S passing through p_1, p_2, p_4, it lies either inside or outside S. By Proposition 4.9, p_3 must lie inside S. Hence, of the two triangulations of R, the one containing $\Delta p_1p_2p_4$ both fails to be the fatter of the two and fails to be the Q–Delaunay triangulation. □

Proof of Theorem 4.8. Let $\mathcal{T}$ be a triangulation of X with optimal angle sequence. In other words $\mathcal{T}$ is the fattest triangulation or possibly tied for fattest. Consider any quadrilateral R in $\mathcal{T}$. We will show that the empty Q–circle condition holds for R, meaning that each of the two triangles of R has the property that its circumscribed Q–circle does not contain the fourth point of R. This is true automatically if R is not convex. So assume R is convex. Then the optimality of $\mathcal{T}$ implies that the triangulation of R also optimizes the angle sequence (six angles) for R. Hence, by Lemma 4.10, the empty Q–circle condition holds for R.

Since each quadrilateral satisfies the empty circle condition, it follows that the polygonal surface obtained by lifting the triangles of $\mathcal{T}$ up to the graph of Q in $\mathbb{R}^3$ is locally convex along all edges. It is therefore a convex surface. It now follows that $\mathcal{T}$ must be the Q–Delaunay triangulation. □

5. Outline of applications

We now briefly outline some possible applications of higher signature Delaunay triangulations.

5.1. **Triangulations in Minkowski space.** The most direct application is to define triangulations with a given vertex sets in spacelike position in Minkowski space, or in other constant curvature Lorentz spaces.

Suppose for instance that $\Sigma \subset \mathbb{R}^{d-1,1}$ is a spacelike hypersurface in the d-dimensional Minkowski space, and let X be a finite set of points of Σ. The Q-Delaunay decomposition with vertex set X — where Q denotes the Minkowski scalar product — provides a triangulation of the convex hull of X which we expect to have good properties with respect to the ambient geometric structure.

5.2. **Analysis of discrete functions.** There are possible applications of the higher signature Delaunay decomposition not only for sets of points in Minkoski space, but also in other situations where one dimension (or more) plays a role fundamentally different from the others.

Consider for instance a finite set $X_0 \in \mathbb{R}^{d-1}$ and a k-Lipschitz function $u : X_0 \to \mathbb{R}$ obtained by sampling a function $v : \mathbb{R}^{d-1} \to \mathbb{R}$ at the points of X_0. After multiplying u by a constant, we can suppose that u is k-Lipschitz, for some $k < 1$. Let $X \subset \mathbb{R}^d$ be the graph of u. We can then consider the Q-Delaunay decomposition with vertex set X, where

$$Q = dx_1^2 + \cdots + dx_{d-1}^2 - dx_d^2 \ .$$

We expect that the combinatorics of $\mathrm{Del}_Q(X)$ can be useful in analysing the points of X_0 where u takes particularly interesting or relevant values. Typically, points of X which are vertices of "large" simplices of $\mathrm{Del}_Q(X)$ could be particularly relevant.

For functions which are not *a priori* Lipschitz, it might be more relevant to consider the Q'-Delaunay decomposition with vertex set X, where Q' is the degenerate quadric

$$Q' = dx_1^2 + \cdots + dx_{d-1}^2 \ .$$

5.3. **Proximity Graphs.** Proximity graphs are undirected graphs in which the points spanning an edge are close in some defined sense.

5.3.1. *Minimum Spanning Tree and Traveling Salesperson Cycle.* Given a finite set $X \subset \mathbb{R}^d$, a *Euclidean Minimum Spanning Tree* $\mathrm{MST}(X)$ of X is an (undirected) graph with minimum summed Euclidean edge lengths that connects all points in X and which has only the points of X as vertices. It is easy to see that it has no cycle (that is, that it is really a tree). We generalize this notion to the setting of points X in spacelike and general position with respect to a quadratic form Q on $\mathbb{R}^d$. We may measure distance between two points $x_1, x_2 \in X$ by the function $d_Q(x_1, x_2) = \sqrt{Q(x_1 - x_2)}$. As in the Euclidean case, we define a *Q-minimum spanning tree* $\mathrm{MST}_Q(X)$ of X as a tree with spacelike edges of minimum total length having as vertices exactly the points of X. Any such tree satisfies the following remarkable property, well-known in the case Q is the Euclidean norm.

Proposition 5.1. *Let X be in spacelike and general position with respect to the quadratic form Q. Then a Q-minimum spanning tree $\mathrm{MST}_Q(X)$ is a sub-graph of the one-skeleton of the Q-Delaunay triangulation $\mathrm{Del}_Q(X)$.*

Proof. We need only to treat the case that Q is non-degenerate. The case of Q degenerate follows from Theorem 3.6 by approximating Q by a one-parameter family of non-degenerate quadratic forms.

Consider an edge $[x, y]$ of $\mathrm{MST}_Q(X)$ spanned by two vertices $x, y \in X$. Let $c \in [x, y]$ be the center point of the interval and consider the Q-ball B centered at c and of radius $d_Q(x, y)/2$. Specifically, B is defined by the inequality $Q(z - c) < Q(y - c) = Q(x - c)$. We show that all other points $z \in X$ lie outside of B. Suppose by contradiction that there is $z \in X$ such that $Q(z - c) < Q(y - c) = Q(x - c)$. We first show that $Q(z - x), Q(z - y) < Q(x - y)$. Indeed, if Q is Euclidean, then this follows easily from the triangle inequality. However, we note that if Q is not positive definite, the triangle inequality fails even for points in spacelike position. Set $D^2 = Q(x - y)$. Then

$$\begin{aligned} D^2 &= \langle x - y, x - y\rangle_Q \\ &= Q(x) + Q(y) + 2\langle x, y\rangle_Q \ , \end{aligned}$$

and therefore

$$2\langle x, y\rangle_Q = -D^2 + Q(x) + Q(y) \ .$$

Next, using that $Q\left(z - \frac{x}{2} - \frac{y}{2}\right) < \left(\frac{D}{2}\right)^2$ and the above expression for $\langle x, y\rangle_Q$, we have:

$$Q(z) + \frac{1}{4}Q(x) + \frac{1}{4}Q(y) - \langle z, x\rangle_Q - \langle z, y\rangle_Q + \frac{1}{2}\langle x, y\rangle_Q < \frac{D^2}{4}$$

so that

$$Q(z) + \frac{1}{2}Q(x) + \frac{1}{2}Q(y) - \langle z, x\rangle_Q - \langle z, y\rangle_Q < \frac{D^2}{2}$$

and thus

$$Q(z - x) + Q(z - y) < D^2 \ .$$

Hence, since both $Q(z-x), Q(z-y)$ are positive, we have that $Q(z-x), Q(z-y) < D^2$ and so $d_Q(x, z), d_Q(y, z) < D = d_Q(x, y)$. Next, there is some path in $\mathrm{MST}_Q(X)$ that connects z to either x or y not passing through $[x, y]$. Without loss in generality, z is connected to y by such a path. Replacing the segment $[x, y]$ with $[z, x]$ yields a spanning tree with smaller total Q-length which contradicts that $\mathrm{MST}_Q(X)$ was minimal. Hence there are no points of X inside the ball B. Hence if X' denotes the lift of X to the graph of Q, as in the proof of Theorem A, then the plane whose intersection with the graph of Q projects to ∂B is a support plane to the bottom

of the convex hull $\mathscr{C}$ of X' which contains the lift $[x', y']$ of the edge $[x, y]$. Hence $[x', y']$ is an edge of the convex hull of X' and so $[x, y]$ is an edge of $\mathrm{Del}_Q(X)$. □

In the Euclidean plane, Proposition 5.1 has been used to provide algorithms calculating $\mathrm{MST}(X)$ in time $O(n \log n)$, where $n = |X|$, while algorithms not using the Delaunay decomposition of X run in $O(n^2)$. In higher dimension, that is, if $d \geq 3$, finding an optimal algorithm remains an open problem. For higher signatures, finding an optimal algorithm to construct a minimum spanning tree for a set X of points in spacelike position appears to be an open question. We did not investigate whether algorithms that apply in the Euclidean case extend to higher signatures. It might be relatively easy in the Minkowski plane, since the spacelike condition is then quite strong, but it is conceivable that Proposition 5.1 can be used for this question in the 3-dimensional Minkowski space.

5.3.2. *Relative neighborhood and Gabriel graph.* There are also natural higher signature analogues of the relative neighborhood graph, introduced and studied by Toussaint [16] in 1980 in the setting of the Euclidean plane, and the Gabriel graph, introduced by Gabriel–Sokal [11] in 1969 in the setting of Euclidean space of any dimension. Let X be a finite set in $\mathbb{R}^d$ which is in spacelike position. Then the *relative neighborhood graph* $\mathrm{RNG}_Q(X)$ of X is the (undirected) graph which has an edge connecting x to y if $d_Q(x, y) < \max\{d_Q(z, x), d_Q(z, y)\}$ for all $z \in X \setminus \{x, y\}$.

Next assume Q is non-degenerate. The *Gabriel graph* $\mathrm{GG}_Q(X)$ of X has an edge $[x, y]$ for $x, y \in X$ if the open Q-ball whose diameter is the segment $[x, y]$ is empty of all other points of X. This notion does not seem to have an analogue for Q degenerate. It is straightforward from the proof of Proposition 5.1 to show:

Proposition 5.2. *In the case Q non-degenerate,*

$$\mathrm{MST}_Q(X) \subset \mathrm{RNG}_Q(X) \subset \mathrm{GG}_Q(X) \subset \mathrm{Del}_Q(X).$$

In the case Q degenerate,

$$\mathrm{MST}_Q(X) \subset \mathrm{RNG}_Q(X) \subset \mathrm{Del}_Q(X).$$

5.4. **Minimizing interpolation error.** The Delaunay triangulation $\mathrm{Del}_Q(X)$ can also be characterized from a functional approximation point of view. For example, in the Euclidean setting, Lambert [13] showed that the classical Delaunay triangulation maximizes the arithmetic mean of the radius of inscribed circles of the triangles, while Rippa [14] proved that it minimizes the integral of the squared gradient.

Let $\Omega \subset \mathbb{R}^d$ be a bounded domain, let T be a triangulation of Ω, and let f be a function defined on Ω. We define

$$\mathcal{Q}(T, f, p) = \|f - \hat{f}_T\|_{Q, L^p(\Omega)} \ ,$$

where $\hat{f}_T$ is the linear interpolation of f based on the triangulation T of Ω, and where we use the quadratic form Q. The following result, due to Chen–Xu [3] in the Euclidean setting, characterizes the Q-Delaunay triangulation as the optimal triangulation for the piecewise linear interpolation over a given point set X of the function $Q(x)$. The proof of Chen–Xu, which uses the convex hull interpretation of the Delaunay triangulation, is easily generalized to the higher signature setting under the additional hypothesis that the points are in spacelike position.

Theorem 5.1. *Let $X \subset \mathbb{R}^d$ be a finite set in spacelike position with respect to Q, then*

$$\mathcal{Q}(\mathrm{Del}_Q(X), Q(\cdot), p) = \min_{T \in \mathcal{P}(X)} \mathcal{Q}(T, Q(\cdot), p), \ \textit{for } 1 \leq p \leq \infty \ ,$$

where $\mathcal{P}(X)$ is the set of all triangulations that have X as vertices and Ω is the convex hull of X.

5.5. **Other applications.** There are many more other well-known and important applications/generalizations of the classical Delaunay triangulation, for example the notion of α–shapes [10] or β–skeletons [12], which have important applications in pattern recognition, digital shape sampling and processing, and structural molecular biology, among others. We encourage the reader to investigate generalizations of such constructions to the higher signature setting as needed.

5.6. **Implementation.** There are several efficient algorithms that can be used to compute the Delaunay triangulation of a set of points in the Euclidean plane or in higher-dimensional space, for instance an incremental flipping algorithm [9] or a divide-and-conquer algorithm [4], both running in $O(n \log(n))$ in the plane (see [1] for efficient implementations). We have not investigated the extent to which these algorithms generalize readily to the setting of Q–Delaunay triangulations, nor have we investigated whether similar efficiency can be achieved. Of course any convex hull algorithm may applied to directly calculate Q–Delaunay decompositions as in the proof of Theorem A.

In an online appendix to this paper[1], we provide a crude implementation of the computation of Delaunay decomposition for the Euclidean, Minkowski and degenerate bilinear forms in $\mathbb{R}^3$. This implementation is provided for illustration only, and is not intended to be computationally efficient.

References

[1] Jean-Daniel Boissonnat, Olivier Devillers, Sylvain Pion, Monique Teillaud, and Mariette Yvinec, *Triangulations in CGAL*, Comput. Geom. **22** (2002), no. 1-3, 5–19, 16th ACM Symposium on Computational Geometry (Hong Kong, 2000). MR 1893651

[2] Kevin Q Brown, *Voronoi diagrams from convex hulls*, Inf. Process. Lett. **9** (1979), no. 5, 223–228.

[3] Long Chen and Jin-chao Xu, *Optimal Delaunay triangulations*, J. Comput. Math. **22** (2004), no. 2, 299–308, Special issue dedicated to the 70th birthday of Professor Zhong-Ci Shi. MR 2058939 (2005c:41042)

[4] Paolo Cignoni, Claudio Montani, and Roberto Scopigno, *Dewall: A fast divide and conquer delaunay triangulation algorithm in e^d*, Computer-Aided Design **30** (1998), no. 5, 333–341.

[5] Daryl Cooper, Jeffrey Danciger, and Anna Wienhard, *Limits of geometries*, preprint arXiv:1408.4109 (2014).

[6] Jeffrey Danciger, *A geometric transition from hyperbolic to anti-de Sitter geometry*, Geom. Topol. **17** (2013), no. 5, 3077–3134. MR 3190306

[7] Jeffrey Danciger, Sara Maloni, and Jean-Marc Schlenker, *Polyhedra inscribed in a quadric*, preprint arXiv:1410.3774 (2014).

[8] Boris Delaunay, *Sur la sphère vide. A la mémoire de George Voronoi.*, Bull. Acad. Sc. USSR, Ph. Ser. **6** (1934), 793–800.

[9] H. Edelsbrunner and N. R. Shah, *Incremental topological flipping works for regular triangulations*, Algorithmica **15** (1996), no. 3, 223–241. MR 1368251 (96j:65161)

[10] Herbert Edelsbrunner, David G. Kirkpatrick, and Raimund Seidel, *On the shape of a set of points in the plane*, IEEE Trans. Inform. Theory **29** (1983), no. 4, 551–559. MR 713690 (84m:52016)

[11] K Ruben Gabriel and Robert R Sokal, *A new statistical approach to geographic variation analysis*, Syst. Zool. **18** (1969), no. 3, 259–278.

[12] David G Kirkpatrick and John D Radke, *A framework for computational morphology*, Computational Geometry, Machine Intelligence and Pattern Recognition 2, 1985.

[13] Timothy Lambert, *The Delaunay triangulation maximizes the mean inradius*, Proc. 6th Canad. Conf. Comput. Geom., 1994, pp. 201–206.

[14] Samuel Rippa, *Minimal roughness property of the Delaunay triangulation*, Comput. Aided Geom. Design **7** (1990), no. 6, 489–497. MR 1079398 (91j:65018)

[15] Igor Rivin, *A characterization of ideal polyhedra in hyperbolic 3-space*, Ann. of Math. (2) **143** (1996), no. 1, 51–70. MR 1370757 (96i:52008)

[16] Godfried T. Toussaint, *The relative neighbourhood graph of a finite planar set*, Pattern Recognition **12** (1980), no. 4, 261–268. MR 591314 (83h:52001)

[1] `http://math.uni.lu/schlenker/programs/highersign/highersign.html`

Department of Mathematics, University of Texas at Austin
E-mail address: jdanciger@math.utexas.edu
URL: www.ma.utexas.edu/users/jdanciger

Department of Mathematics, Brown University
E-mail address: sara_maloni@brown.edu
URL: http://www.math.brown.edu/∼maloni

Department of mathematics, University of Luxembourg
E-mail address: jean-marc.schlenker@uni.lu
URL: math.uni.lu/schlenker